Non-Interactive Zero-Knowledge Arguments with Certified Deletion*

Kasra Abbaszadeh[1] and Jonathan Katz[2]

[1] University of Maryland
kasraz@umd.edu
[2] Google
jkatz2@gmail.com

Abstract. We introduce the notion of non-interactive zero-knowledge (NIZK) arguments with *certified deletion*, a new primitive that enables the recipient of a (quantum) NIZK argument to delete it and obtain a (classical) certificate proving such deletion. We formalize this notion and propose two candidate constructions from standard cryptographic assumptions. Our first construction is based on classical NIZK arguments and quantum-hard one-way functions, but requires both the prover and verifier to run quantum algorithms. We then present an extension based on the learning with errors problem that allows the prover to be classical.

Our results have applications to signatures of knowledge and anonymous credentials with certified deletion, which we also define and construct.

Keywords: quantum cryptography, certified deletion, zero-knowledge arguments, signatures of knowledge, anonymous credentials.

1 Introduction

Quantum computation enables novel cryptographic capabilities that are known to be impossible in the classical world. These include primitives such as quantum money [46], unconditionally secure key agreement [14], one-shot signatures [6], copy protection [1], unclonable encryption [25,18], and certified deletion [17,10].

Quantum copy protection [1] ensures that, given a single copy of a function f, no adversary can generate two copies of f. Although this has many applications (e.g., public-key quantum money [2]), building copy protection from standard assumptions remains a challenging open problem. Motivated by this, Ananth et al. [7] introduced the weaker notion of *secure software leasing*. At a high level, this notion enables a function f to be encoded in a quantum state that can be given to a user; later, the user can generate a deletion certificate proving that it deleted its copy, at which point the user is no longer able to execute f. Roughly speaking, then, copy protection prevents cloning altogether, whereas software leasing prevents cloning only for users who produce a valid deletion certificate.

* Work supported in part by NSF award CNS-2154705.

Several prior works studied software leasing—also known as revocation—for specific cryptographic functionalities. Ananth et al. [8] and Agrawal et al. [5] independently introduced revocable public-key encryption, where a decryption key is leased to a user; when the key is returned, the user loses the ability to decrypt. Morimae et al. [36] studied digital signature schemes where either the signing key or signatures can be revoked. We extend this line of research and consider revocable non-interactive zero-knowledge (NIZK) arguments [15]. An NIZK argument system enables a prover holding a witness for an NP statement to generate a proof of the statement without revealing any information about the witness. A classical NIZK argument can be verified arbitrarily many times by anyone who obtains it. In some settings, however, a prover may wish to allow verification only for a limited time. To address this, we introduce *NIZK with certified deletion* (NIZK-CD), which enables recipients to certifiably delete arguments they are given. The certificate can be validated by the prover; once a recipient returns such a certificate, it can no longer output a valid argument.

1.1 Our Contributions

We define NIZK-CD and propose constructions in the common reference string (CRS) model based on quantum-secure NIZK arguments and one-way functions. Deletion certificates and their corresponding validation algorithm are classical. In our first construction, however, the prover and verifier are required to run quantum algorithms. We then extend this construction, based on the learning with errors (LWE) assumption [42], to allow the prover to be implemented by a classical algorithm that remotely prepares the proof for a (quantum) verifier.

As a natural application, we construct *signatures of knowledge* with certified deletion from NIZK-CD and use them to construct *anonymous credentials* with certified deletion, without blocklisting or expiration-based techniques [13,4]. Goyal et al. [26] and Jawale et al. [30] previously built (anonymous) revocable credentials from unclonable NIZKs. Unlike ours, their schemes are based on iO and require quantum communication for credential generation and deletion.

1.2 Related Work

We discuss prior and concurrent work on copy protection and certified deletion.

Copy protection and unclonable primitives. Aaronson [1] defined copy protection and presented a copy-protection scheme for any unlearnable Boolean function relative to a quantum oracle. Later [3], it was shown that a classical oracle suffices. Coladangelo et al. [22] proposed copy protection for (multi-bit) point and compute-and-compare functions in the quantum random-oracle model. Copy protection for decryption schemes and pseudorandom functions can be realized from compute-and-compare obfuscation for unpredictable distributions, indistinguishability obfuscation (iO), and one-way functions [5]. Liu et al. [34] built bounded collusion-resistant copy protection for functionalities including pseudorandom functions based on iO and the LWE assumption.

Recently, Goyal et al. [26], and Jawale et al. [30] studied copy protection for NIZK arguments, also called *unclonable NIZK*. They showed that this notion is equivalent to public-key quantum money [2], which is currently only known to exist under the assumption of iO.

Certified deletion and revocable cryptography. Unruh [45] introduced the notion of revocable encryption where the recipient of a quantum ciphertext can delete the ciphertext and lose all information about the message. Broadbent et al. [17] studied certified deletion for quantum one-time pad encryption. There have been several subsequent works [28,40,10,9,11,20] extending these ideas to other primitives such as public-key and attribute-based encryption. Kitagawa et al. [33] and Bartusek et al. [11] replaced privately verifiable certificates with publicly verifiable ones under the assumption of one-way functions and one-way state generators [37]. Certified deletion has also been studied for the revocation of cryptographic keys [5,8,20], digital signatures [36], secret sharing [12,31], and obfuscation [9]. We remark that revocation has also been considered for general programs by Ananth et al. [7] based on quantum-secure iO.

Concurrent work. Concurrent work by Çakan et al. [19] defines NIZK with *certified deniability*, a stronger guarantee than certified deletion. Roughly, NIZK with certified deletion guarantees that once a user certifiably deletes a proof, the user cannot generate a proof that will be accepted by the honest verification algorithm; NIZK with certified deniability ensures that once a user certifiably deletes a proof, they have no advantage in convincing even a dishonest verifier (who uses an arbitrary verification algorithm) about the truth of the statement. Çakan et al. [19] proposed a construction of NIZK with certified deniability in the random-oracle model. Their construction has quantum deletion certificates; that is, certificate validation in their scheme must be done by a quantum algorithm. An open problem is to realize NIZKs with certified deniability where the deletion certificate is classical. Çakan et al. [19] further showed the impossibility of NIZK with certified deniability in the plain model under black-box security reductions. We believe the same result holds in the presence of a classical common reference string (CRS). Another open problem is to bypass this impossibility result.

2 Technical Overview

In this section, we give a high-level overview of our techniques.

2.1 Definition of NIZK-CD

An NIZK-CD scheme consists of algorithms $\langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}, \mathsf{Delete}, \mathsf{Certify} \rangle$. The first three algorithms (Setup, Prove, and Verify) are defined largely as for standard NIZK. In particular, Setup generates a CRS and its trapdoor, and Prove and Verify are used, respectively, to generate and verify a (quantum) proof π. The only difference is that Prove additionally outputs a (classical) certification

key ck that is kept private. The Delete algorithm collapses π into a (classical) deletion certificate cert that can be validated by running $\mathsf{Certify}$ on cert and ck.

An NIZK-CD scheme satisfies the standard requirements of NIZK, i.e., completeness, soundness, and zero knowledge. For deletion security, one natural definition is as follows. Let $\mathcal{D}$ be an efficiently sampleable quantum-hard distribution over instance-witness pairs. Then we could require that if we sample $(x, w) \leftarrow \mathcal{D}$ and give an adversary x and an NIZK-CD proof π generated using x, w, it should be infeasible for the adversary to output both a valid deletion certificate cert as well as a valid proof π^* for x. This definition, however, does not capture an adversary who receives more than one proof, and it does not offer any non-malleability guarantee. To address these limitations, we strengthen the definition to require that for any adversary given oracle access to the prover algorithm, if the adversary outputs both a valid deletion certificate and an accepting proof, then it is possible to extract a corresponding witness.

2.2 Constructions of NIZK-CD

We propose a generic compiler that transforms any NIZK to a NIZK-CD using a non-interactive commitment scheme with statistical binding and computational hiding (against quantum algorithms); such commitments can be realized from quantum-hard one-way functions [23] in the CRS model. Let $c \leftarrow \mathsf{Com}(\mathsf{crs}, m, r)$ denote a commitment to a message m under randomness r using CRS crs.

At a high level, our construction proceeds as follows. Given an NP relation R, an NIZK-CD proof for an instance-witness pair $(x, w) \in R$ is generated by having the prover sample r_0 and r_1 and then generate $c_b = \mathsf{Com}(\mathsf{crs}, b, r_b)$ for $b \in \{0, 1\}$. The prover then produces an NIZK proof σ that

$$\exists w, r_0, r_1 : \ (x, w) \in R \bigvee \left(\bigwedge_{b \in \{0,1\}} c_b = \mathsf{Com}(\mathsf{crs}, 1-b, r_b) \right).$$

Assuming σ is valid, such a proof is convincing if it is possible to check that some c_b is a commitment to b (rather than $1 - b$). To achieve the latter, the prover includes the quantum state $|\psi\rangle \propto |0\rangle|r_0\rangle + (-1)^u|1\rangle|r_1\rangle$ (for uniform u) as part of the proof, while retaining $\mathsf{ck} := (r_0, r_1, u)$ as the certification key.

A proof $(|\psi\rangle, \sigma, \{c_b\}_{b\in\{0,1\}})$ is verified as follows. Let $U^{\mathsf{com}}_{c_0,c_1}$ be the unitary defined by

$$|b\rangle|r\rangle|0\rangle \xrightarrow{U^{\mathsf{com}}_{c_0,c_1}} |b\rangle|r\rangle|\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}, b, r, c_b)\rangle,$$

where $\mathsf{Cmt\text{-}Cmp}$ is a commit-and-compare function that commits to b using randomness r and returns 1 iff the result is equal to the hardcoded string c_b. The verifier checks σ and then applies $U^{\mathsf{com}}_{c_0,c_1}$ to $|\psi\rangle|0\rangle$, measures the third register, and verifies that the result is 1. For honestly generated proofs, the result is always 1. Conversely, if the result is 1 then at least one of the commitments c_b is a commitment to b and so x must be in the language. Note also that for honestly generated proofs the post-measurement state is unchanged; therefore, the proof can be re-verified multiple times.

The verifier can delete $|\psi\rangle$ by applying a Hadamard measurement, which yields a string d with $d \cdot (r_0 \oplus r_1) = u$; note that d can be verified using ck. This measurement causes all other information about r_0, r_1 to be lost as the state collapses to $|0\rangle + (-1)^{d\cdot(r_0 \oplus r_1)}|1\rangle$. Formally, deletion security follows from the adaptive hardcore-bit property of one-way functions [11,36]. This property states that for any one-way function f, no efficient adversary given $y_0 = f(r_0)$, $y_1 = f(r_1)$, and a superposition $|r_0\rangle + (-1)^u|r_1\rangle$ can output both a preimage r such that $f(r) \in \{y_0, y_1\}$ and a string d such that $d \cdot (r_0 \oplus r_1) = u$ with advantage significantly better than 1/2. (The advantage can be reduced to negligible via parallel repetition.) Since the commitment is hiding and binding, it can be viewed as a one-way function for which the adaptive hardcore-bit property holds.

Classical prover. We observe that the prover only needs quantumness to create the state $|\psi\rangle$ and send it to the verifier. Other components of proof generation and deletion require only classical computation and communication. We present an extension of our construction that allows for remote preparation of $|\psi\rangle$ at the verifier; this can be done before x, w are known.

We first recall the notion of *noisy trapdoor claw-free (NTCF)* functions. For a key k and a pair of functions $f_{k,0}, f_{k,1} : \mathcal{X} \to \mathcal{Y}$ of an NTCF family, it is computationally hard to find a *claw*, i.e., a pair (r_0, r_1) where $f_{k,0}(r_0) = f_{k,1}(r_1)$. On the other hand, there is a trapdoor td that allows for efficient computation of preimages r_0 and r_1 of any image point y, i.e., such that $f_{k,b}(r_b) = y$ for $b \in \{0, 1\}$. Moreover, the adaptive hardcore-bit property holds, i.e., given y and $|r_0\rangle + |r_1\rangle$, where r_0 and r_1 are preimages of y, no adversary can return one of the preimages along with a string d such that $d \cdot (r_0 \oplus r_1) = 0$. An NTCF family can be constructed under the LWE assumption [16].

We now define the following protocol:

1. The prover generates (k, td) for an NTCF family and sends k to the verifier.
2. The verifier generates $|\phi\rangle := \sum_{r\in\mathcal{X}} |0\rangle|r\rangle|0\rangle + |1\rangle|r\rangle|0\rangle$. Define $U_{f_{k,0},f_{k,1}}$ via
$$|b\rangle|r\rangle|0\rangle \xrightarrow{U_{f_{k,0},f_{k,1}}} |b\rangle|r\rangle|f_{k,b}(r)\rangle.$$
The verifier applies $U_{f_{k,0},f_{k,1}}$ to $|\phi\rangle$ and measures the third register to obtain an image y and a state $|\psi\rangle = |0\rangle|r_0\rangle + |1\rangle|r_1\rangle$. It sends y to the prover.
3. The prover uses td to recover r_0, r_1 from y, and generates $\sigma, \{c_b\}$ as before.

Although the state-preparation phase is interactive, it is independent of x, w (i.e., it is *instance-independent*); moreover, the final proof $(|\psi\rangle, \sigma, \{c_b\})$ can be verified by anyone. We stress in particular, that soundness holds even if $|\psi\rangle$ is generated by a malicious prover, rather than via the process above. Thus, the prover need not perform state preparation with the eventual verifier, but could instead rely on any untrusted third party.

3 Preliminaries

We use λ to denote the security parameter. We use negl as a generic negligible function. For a set S, we use $x \leftarrow S$ to indicate that x is sampled uniformly

from S. We define $[n] := \{0, 1, \ldots, n-1\}$. The term PPT stands for probabilistic polynomial time and QPT stands for quantum polynomial time.

Quantum conventions. A register X is a Hilbert space $\mathbb{C}^{2^n}$. An n-qubit pure state on register X is a unit vector $|\psi\rangle \in \mathbb{C}^{2^n}$. A mixed state on register X, described by a density matrix $\rho \in \mathbb{C}^{2^n \times 2^n}$, is a positive semi-definite Hermitian operator with trace 1. Also, a quantum operation F is a completely positive trace-preserving map from a register X to a register Y, i.e., on input a density matrix ρ on register X, the operation F returns $F(\rho)$ on register Y. A unitary operation $U : X \to X$ is a quantum operation that satisfies $U^\dagger U = UU^\dagger = I^X$, where I^X is identity. A projector Π is a Hermitian operator such that $\Pi^2 = \Pi$. A projective measurement is a collection of projectors $\{\Pi_i\}_i$ with $\sum_i \Pi_i = I$.

Densities and distances. Let $\mathcal{X}$ be a finite domain. A density on $\mathcal{X}$ is a function $f : \mathcal{X} \to [0, 1]$ such that $\sum_{x \in \mathcal{X}} f(x) = 1$. $\mathcal{D}_\mathcal{X}$ denotes the set of densities on $\mathcal{X}$. For any $f \in \mathcal{D}_\mathcal{X}$, $\mathsf{Supp}(f) = \{x \in \mathcal{X} : f(x) > 0\}$. Given two densities f_0 and f_1 on $\mathcal{X}$, the Hellinger distance is defined as

$$H^2(f_0, f_1) := 1 - \sum_{x \in \mathcal{X}} \sqrt{f_0(x) f_1(x)}.$$

For two density matrices ρ and σ, their trace distance is defined as

$$\|\rho - \sigma\|_{\mathsf{tr}} = \frac{1}{2}\|\rho - \sigma\|_1 = \frac{1}{2}\mathsf{Tr}\left[\sqrt{(\rho - \sigma)^2}\right],$$

where $\|\cdot\|_1$ is the trace norm.

Lemma 3.1. *Let $\mathcal{X}$ be a finite set, $f_0, f_1 \in \mathcal{D}_\mathcal{X}$, and $|\psi_b\rangle := \sum_{x \in \mathcal{X}} \sqrt{f_b(x)}\,|x\rangle$ for $b \in \{0, 1\}$. We have*

$$\||\psi_0\rangle\langle\psi_0| - |\psi_1\rangle\langle\psi_1|\|_{\mathsf{tr}} = \sqrt{1 - (1 - H^2(f_0, f_1))^2}.$$

Theorem 3.1 (Holevo-Helstrom [27,29]). *Consider the experiment in which one of two quantum states ρ and σ is given to some distinguisher, each with probability $\frac{1}{2}$. The advantage of any distinguisher in correctly determining which state was sent is at most $\frac{1}{2} + \frac{1}{2}\|\rho - \sigma\|_{\mathsf{tr}}$.*

3.1 The Learning with Errors Assumption

We recall the definition of the learning with errors (LWE) problem. For real positive B and integer q, the discrete truncated Gaussian distribution over $\mathbb{Z}_q$ with parameter B is a distribution on $\{x \in \mathbb{Z}_q : \|x\| \leq B\}$ with density

$$D_{\mathbb{Z}_q, B}(x) := \frac{e^{\frac{-\pi\|x\|^2}{B}}}{\sum_{x' \in \mathbb{Z}_q, \|x'\| \leq B} e^{\frac{-\pi\|x'\|^2}{B}}}.$$

For $d > 1$, the truncated discrete Gaussian distribution over $\mathbb{Z}_q^d$ with parameter B is a distribution on $\{x \in \mathbb{Z}_q^d : \|x\| \leq B\sqrt{d}\}$ with density

$$\forall x = (x_1, x_2, \ldots, x_d) \in \mathbb{Z}_q^d : \quad D_{\mathbb{Z}_q^d, B}(x) = D_{\mathbb{Z}_q, B}(x_1) D_{\mathbb{Z}_q, B}(x_2) \cdots D_{\mathbb{Z}_q, B}(x_d).$$

We now define the LWE problem.

Definition 3.1 (LWE). *Let $n(\lambda), m(\lambda), q(\lambda)$ be polynomials in λ. Moreover, let $\mathcal{X} = \mathcal{X}(\lambda)$ be a distribution over $\mathbb{Z}$. The $\mathsf{LWE}_{m,n,q,\mathcal{X}}$ problem is to distinguish between the distributions $(A, As + e)$ and (A, u), where $A \leftarrow \mathbb{Z}_q^{m \times n}$, $s \leftarrow \mathbb{Z}_q^n$, $e \leftarrow \mathcal{X}^m$, $u \leftarrow \mathbb{Z}_q^m$.*

We assume that no QPT algorithm can solve $\mathsf{LWE}_{m,n,q,\mathcal{X}}$ with non-negligible advantage, even given access to a quantum polynomial-size advice depending on the parameters m, n, q, and $\mathcal{X}$ of the problem. We refer to this as the $\mathsf{LWE}_{m,n,q,\mathcal{X}}$ assumption. It can be shown [42,38] that for any $\alpha > 0$ such that $\sigma = \alpha q \geq 2\sqrt{n}$, the $\mathsf{LWE}_{m,n,q,D_{\mathbb{Z}_{q,\sigma}}}$ problem is at least as hard as solving the shortest independent vector problem on n-dimensional lattices within a factor of $\gamma = \tilde{O}(n/\alpha)$ in the *worst case*, where $\tilde{O}$ hides logarithmic factors. The best known algorithm to solve the problem runs in time $2^{\tilde{O}(n/\log \gamma)}$. We assume hardness against QPT adversaries for γ super-polynomial in n.

3.2 Noisy Trapdoor Claw-Free Hash Function Families

We recall noisy trapdoor claw-free (NTCF) hash functions [16]. Given two finite sets $\mathcal{X}$ and $\mathcal{Y}$, a trapdoor claw-free family of functions satisfies the following properties. For a public key k, the functions $\{f_{k,b} : \mathcal{X} \to \mathcal{Y}\}_{b \in \{0,1\}}$ are both injective and have the same range and are invertible given a trapdoor td, i.e., on input td and an image $y \in \mathcal{Y}$ it is feasible to efficiently output $x_0, x_1 \in \mathcal{X}$ such that $f_0(x_0) = f_1(x_1) = y$. Furthermore, the pair of functions should be claw-free, i.e., it is computationally hard for an attacker to find two preimages x_0, x_1 such that $f_0(x_0) = f_1(x_1)$ without the trapdoor. The functions also satisfy the adaptive hardcore bit property, which states that it is computationally hard for an attacker to output a non-trivial (b, d, x_b) with a probability non-negligibly greater than $\frac{1}{2}$ the such that the equation $d \cdot (x_0 \oplus x_1) = 0$ is satisfied; note that x_{1-b} is a unique element such that $f_{1-b}(x_{1-b}) = f_b(x_b)$.

Unfortunately, we are not aware of any exact constructions of the trapdoor claw-free functions. Instead, Brakerski et al. [16] proposed a construction for noisy trapdoor claw-free functions, which relaxes the requirements as follows. First, the range of functions is not $\mathcal{Y}$, but instead $\mathcal{D}_\mathcal{Y}$, the set of probability densities over $\mathcal{Y}$. The trapdoor injective pair property is then defined according to the support of the output densities. The use of densities as the output of the functions requires considering additional requirements. In this paper, we need a quantum polynomial-time algorithm that efficiently prepares a superposition over the range of the function, i.e., given a function key k and a bit $b \in \{0, 1\}$,

the algorithm can prepare the following quantum state:

$$\frac{1}{\sqrt{\mathcal{X}}} \sum_{x\in\mathcal{X},y\in\mathcal{Y}} \sqrt{f_{k,b}(x)y}|x\rangle|y\rangle.$$

The construction proposed in [16] is unable to exactly produce the above state; however, it is possible to create a superposition with coefficients such that $f_{k,b}(x)$ is approximated by another function $f'_{k,b}(x)$ and the resulting state is within a negligible distance from the above state. $f'_{k,b}(x)$ supports membership checks efficiently without the need for the trapdoor, and the inversion algorithm should operate properly on the images in the support of $f'_{k,b}(x)$. The adaptive hardcore bit property needs to also be slightly modified. The set $\mathcal{X}$ might not be a subset of binary strings. We first assume the existence of an injective, efficiently invertible map $J : \mathcal{X} \to \{0,1\}^w$ and consider the adaptive hardcore bit property to hold for a subset of all non-zero strings. The formal definition is as follows.

Definition 3.2 (NTCF [16]). *$\mathcal{F} := \{f_{k,b} : \mathcal{X} \to \mathcal{D}_\mathcal{Y}\}_{k\in\mathcal{K},b\in\{0,1\}}$ is a NTCF hash function family if the following properties are satisfied.*

Key generation: *$\mathsf{NTCF.Gen}_\mathcal{F}$ samples a key $k \in \mathcal{K}$ and a trapdoor td.*

Trapdoor Injective Pair: *For all $k \in \mathcal{K}$, $b \in \{0,1\}$, distinct $x, x' \in \mathcal{X}$, $\mathsf{Supp}(f_{k,b}(x)) \cap \mathsf{Supp}(f_{k,b}(x')) = \emptyset$. An efficient deterministic algorithm $\mathsf{Inv}_\mathcal{F}$ exists such that for all $k \in \mathcal{K}$, $b \in \{0,1\}$, $x \in \mathcal{X}$, and $y \in \mathsf{Supp}(f_{k,b}(x))$, it holds that $\mathsf{Inv}(td, b, y) = x$. Moreover, given a key $k \in \mathcal{K}$, there exists a perfect matching $\mathcal{R}_k \subseteq \mathcal{X} \times \mathcal{X}$ such that $f_{k,0}(x_0) = f_{k,1}(x_1)$ if and only if $(x_0, x_1) \in \mathcal{R}_k$.*

Range Superposition: *For all $k \in \mathcal{K}$, $b \in \{0,1\}$, $f'_{k,b} : \mathcal{X} \to \mathcal{D}_\mathcal{Y}$ exists where:*

- *For any claw $(x_0, x_1) \in \mathcal{R}_k$ with image $y \in \mathsf{Supp}(f'_{k,b}(x_b))$ it holds that $\mathsf{Inv}_\mathcal{F}(\mathsf{td}, b, y) = x_b$ as well as $\mathsf{Inv}_\mathcal{F}(\mathsf{td}, b \oplus 1, y) = x_{b\oplus 1}$.*
- *There exists an efficient deterministic algorithm $\mathsf{Chk}_\mathcal{F}$ where on input $k \in \mathcal{K}$, $b \in \{0,1\}$, $x \in \mathcal{X}$, $y \in \mathcal{Y}$, it outputs 1 if $y \in \mathsf{Supp}(f'_{k,b}(x))$ and 0 otherwise.*
- *For all $k \in \mathcal{K}$, $b \in \{0,1\}$, we have $\mathbb{E}_{x\leftarrow X}[H^2(f_{k,b}(x), f'_{k,b}(x))] \leq \mathsf{negl}(\lambda)$. In addition, there exists a QPT algorithm $\mathsf{Samp}_\mathcal{F}$ such that on input $k \in \mathcal{K}$ and $b \in \{0,1\}$ outputs the following quantum state.*

$$|\psi'\rangle = \frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x\in\mathcal{X},y\in\mathcal{Y}} \sqrt{(f'_{k,b}(x))(y)}|x\rangle|y\rangle.$$

Given $|\psi\rangle = \frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x\in\mathcal{X},y\in\mathcal{Y}} \sqrt{(f_{k,b}(x))(y)}|x\rangle|y\rangle$, Lemma 3.1 implies that

$$\||\psi\rangle\langle\psi| - |\psi'\rangle\langle\psi'|\|_{\mathsf{tr}} \leq \mathsf{negl}(\lambda).$$

Adaptive Hardcore Bit: *For all keys $k \in \mathcal{K}$ and a polynomial $\ell(\lambda)$:*

- *For all $b \in \{0,1\}$, $x \in \mathcal{X}$, there exists a subset of strings $\mathcal{G}_{k,b,x} \subseteq \{0,1\}^{\ell(\lambda)}$ such that $\Pr_{d\leftarrow\{0,1\}^{\ell(\lambda)}}[d \notin \mathcal{G}_{k,b,x}] \leq \mathsf{negl}(\lambda)$. Moreover, there exists a PPT algorithm to efficiently check membership in $\mathcal{G}_{k,b,x}$ given k, b, x and td.*

- *There exists an efficiently computable injection $J : \mathcal{X} \to \{0,1\}^{\ell(\lambda)}$ such that J can also be efficiently inverted on its range, and the following holds. Let*

  $H_k := \{(b, x_b, d, d \cdot (J(x_0) \oplus J(x_1)) | b \in \{0,1\}, (x_0, x_1) \in R_k, d \in \mathcal{G}_{k,0,x} \cap \mathcal{G}_{k,1,x}\}$,
  $\bar{H}_k := \{(b, x_b, d, u \oplus 1) | (b, x_b, d, u) \in H_k\}$.

  For any QPT *algorithm $\mathcal{A}$ it holds that*

$$\left| \Pr_{\mathsf{td},k \leftarrow \mathsf{NTCF.Gen}(1^\lambda)} [\mathcal{A}(k) \in H_k] - \Pr_{\mathsf{td},k \leftarrow \mathsf{NTCF.Gen}(1^\lambda)} [\mathcal{A}(k) \in \bar{H}_k] \right| \leq \mathsf{negl}(\lambda).$$

Theorem 3.2 ([16]). *Assuming the polynomial-time quantum hardness of LWE, there is an* NTCF *hash function family.*

One can consider an amplified adaptive hardcore bit property such that the adversary cannot return a set $\{(b_i, x_{i,b_i}, d_i, u_i)\}_{i \in [n]}$ where n is polynomial in the security parameter λ, and each tuple satisfies $d_i \cdot (x_{i,b_i} \oplus x_{i,1-b_i}) = u_i$, so that $(x_{i,b_i}, x_{i,1-b_i})$ is a claw. The property is formally defined as follows.

Definition 3.3 (Amplified Adaptive Hardcore Bit). *An* NTCF *function family $\mathcal{F}$ satisfies the amplified adaptive hardcore bit property if, for any* QPT *algorithm $\mathcal{A}$ and a polynomial $n = \ell(\lambda)$, the following is negligible:*

$$\Pr\left[\begin{array}{c} \forall i \in [n] : (k_i, \mathsf{td}_i) \leftarrow \mathsf{NTCF.Gen}(1^\lambda) \\ \{(b_i, x_{i,b_i}, d_i, u_i)\}_{i \in [n]} \leftarrow \mathcal{A}(\{k_i\}_{i \in [n]}) \\ \forall \beta \in \{0,1\} : x_{i,\beta} = \mathsf{Inv}(\mathsf{td}_i, \beta, y_i) \end{array} : \begin{array}{c} \forall i \in [n] : d_i \in \mathcal{G}_{k,0,x} \cap \mathcal{G}_{k,1,x} \\ \wedge \\ u_i = d_i \cdot (J(x_{i,0}) \oplus J(x_{i,1})) \end{array}\right]$$

Theorem 3.3. *[41,32] Any* NTCF *family of functions satisfies the amplified adaptive hardcore bit property.*

Morimae et al. [36] introduced a similar notion for general one-way functions (OWFs) and showed the following result.

Theorem 3.4. *Given $\ell(\lambda), \kappa(\lambda), n(\lambda) \in \mathbb{N}$, for any quantum-hard OWF $f : \{0,1\}^{\ell(\lambda)} \to \{0,1\}^{\kappa(\lambda)}$, and any* QPT *adversary $\mathcal{A}$,*

$$\Pr\left[\begin{array}{c} \forall i \in [n] : x_{i,0}, x_{i,1} \leftarrow \{0,1\}^{2\ell(\lambda)}, u_i \leftarrow \{0,1\} \\ \{(x_i, d_i)\}_{i \in [n]} \leftarrow \mathcal{A}(\otimes_{i \in [n]} \frac{|x_{i,0}\rangle + (-1)^{u_i} |x_{i,1}\rangle}{\sqrt{2}}, \{f(x_{i,b})\}_{i,b}) \end{array} : \begin{array}{c} \wedge_{i \in [n]} f(x_i) \in \{f(x_{i,0}), f(x_{i,1})\} \\ \wedge \\ \wedge_{i \in [n]} d_i \cdot (x_{i,0} \oplus x_{i,1}) = u_i \end{array}\right] \leq \mathsf{negl}(\lambda).$$

3.3 Other Useful Cryptographic Primitives and Lemmas

We recall the notion of commitment schemes in the CRS model, which is a central building block for our NIZK-CD constructions throughout this paper.

Definition 3.4 (Commitment). *Let $n(\lambda), \ell(\lambda), \kappa(\lambda)$ be polynomials.* Com *is a quantum-secure commitment if the following syntax and properties hold.*

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$*: on input λ, output* crs*, and trapdoor* td*.*
- $c \leftarrow \mathsf{Com}(\mathsf{crs}, m, r)$*: on input* crs*, a message $m \in \{0,1\}^{n(\lambda)}$, and randomness $r \in \{0,1\}^{\ell(\lambda)}$, output a commitment $c \in \{0,1\}^{\kappa(\lambda)}$.*

Statistical Binding: *For any algorithm $\mathcal{A}$,*

$$\Pr\left[\begin{array}{c}(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) \\ (m_0, r_0, m_1, r_1) \leftarrow \mathcal{A}(\mathsf{crs})\end{array} : \begin{array}{c}\mathsf{Com}(\mathsf{crs}, m_0, r_0) = \mathsf{Com}(\mathsf{crs}, m_1, r_1) \\ \wedge\ m_0 \neq m_1\end{array}\right] \leq \mathsf{negl}(\lambda).$$

Computational Hiding: *For any* QPT *distinguisher $\mathcal{D}$ and m_0, m_1,*

$$\left|\Pr\left[\begin{array}{l}(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) \\ \quad r_0 \leftarrow \{0,1\}^{\ell(\lambda)} \\ c_0 \leftarrow \mathsf{Com}(\mathsf{crs}, m_0, r_0)\end{array} : \mathcal{D}(\mathsf{crs}, c_0) = 1\right]\right.$$
$$\left. - \Pr\left[\begin{array}{l}(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) \\ \quad r_1 \leftarrow \{0,1\}^{\ell(\lambda)} \\ c_1 \leftarrow \mathsf{Com}(\mathsf{crs}, m_1, r_1)\end{array} : \mathcal{D}(\mathsf{crs}, c_1) = 1\right]\right| \leq \mathsf{negl}(\lambda).$$

Theorem 3.5 ([23]). *Assuming a quantum-hard one-way function, there is a statistically binding, computationally hiding commitment scheme in the CRS model.*

We recall the notion of non-interactive zero-knowledge (NIZK) arguments in the CRS model. We provide two definitions of NIZK as follows. The first one is a single-theorem definition and ensures that no efficient adversary can output an accepting proof for any unsatisfied instance. The second one is a multi-theorem definition and ensures that from any efficient adversary that can output an accepting proof, one can efficiently extract a valid witness to the relation.

Definition 3.5 (NIZK for NP). *Let R be an* NP *relation with a corresponding language L. $\Pi = \langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify} \rangle$ is a quantum-secure NIZK for* NP *in the CRS model if it satisfies the following syntax and security properties.*

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$ *: on input λ, output* crs *and a trapdoor* td*.*
- $\pi \leftarrow \mathsf{Prove}(\mathsf{crs}, x, w)$ *: on input* crs *and a pair $(x, w) \in R$, output a proof π.*
- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, \pi)$ *: on input* crs*, x, and π, output accept or reject.*

Perfect Completeness: *For every security parameter $\lambda \in \mathbb{N}$ and $(x, w) \in R$,*

$$\Pr\left[\begin{array}{l}(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi \leftarrow \mathsf{Prove}(\mathsf{crs}, x, w)\end{array} : \mathsf{Verify}(\mathsf{crs}, x, \pi)\right] = 1.$$

Adaptive Computational Soundness: *For any* QPT *adversary algorithm* $\mathcal{A}$

$$\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}(\mathsf{crs})\end{array}:\mathsf{Verify}(\mathsf{crs},x,\pi)=1\wedge x\notin L\right]\leq\mathsf{negl}(\lambda).$$

Adaptive Computational Zero-Knowledge: *There exists a* QPT *simulator* $\mathsf{Sim}=(\mathsf{Sim}_0,\mathsf{Sim}_1)$ *where for every* QPT *adversary* $\mathcal{A}$ *and distinguisher* $\mathcal{D}$*,*

$$\left|\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\(x,w,\xi)\leftarrow\mathcal{A}(\mathsf{crs})\\\pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\end{array}:\mathcal{D}(\mathsf{crs},x,\pi,\xi)=1\right]-\right.$$
$$\left.\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,w,\xi)\leftarrow\mathcal{A}(\mathsf{crs})\\\pi\leftarrow\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x)\end{array}:\mathcal{D}(\mathsf{crs},x,\pi,\xi)=1\right]\right|\leq\mathsf{negl}(\lambda).$$

Theorem 3.6 ([39]). *Assuming quantum hardness of the LWE assumption, there is a quantum-secure non-interactive, adaptively computationally sound, and adaptively computationally zero-knowledge argument for NP.*

Definition 3.6 (Simulation-Extractable NIZK for NP). *Let* R *be an* NP *relation for language* L*. Then,* $\Pi=\langle\mathsf{Setup},\mathsf{Prove},\mathsf{Verify}\rangle$ *is a quantum-secure simulation-extractable NIZK for NP in the CRS model if it satisfies the properties of Definition 3.5 and also the following additional properties.*

Adaptive Multi-Theorem Computational Zero-Knowledge: *There is a* QPT *simulator* $\mathsf{Sim}=(\mathsf{Sim}_0,\mathsf{Sim}_1)$ *such that for any* QPT *adversary* $\mathcal{A}$*,*

$$\left|\Pr\left[(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda):\mathcal{A}^{\mathsf{Prove}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs})=1\right]-\right.$$
$$\left.\Pr\left[(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda):\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})=1\right]\right|\leq\mathsf{negl}(\lambda),$$

where the simulator algorithm Sim_1 *discards its fourth input value.*

Simulation Soundness: *Let* $\mathsf{Sim}=(\mathsf{Sim}_0,\mathsf{Sim}_1)$ *be the simulator given by the adaptive multi-theorem computational zero-knowledge. For every* QPT *algorithm* $\mathcal{A}$*, the following is negligible:*

$$\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})\end{array}:\mathsf{Verify}(\mathsf{crs},x,\pi)=1\wedge x\notin L\wedge x\notin Q\right],$$

where Q *is the list of queries from* $\mathcal{A}$ *to* Sim_1*.*

Simulation Extractability: *Let* $\mathsf{Sim}=(\mathsf{Sim}_0,\mathsf{Sim}_1)$ *be the simulator given by adaptive multi-theorem computational zero knowledge. There is a* QPT *extractor* Ext *such that for every* QPT *algorithm* $\mathcal{A}$ *the following is negligible:*

$$\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})\\b\leftarrow\mathsf{Verify}(\mathsf{crs},x,\pi)\\w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)\end{array}:b=1\wedge(x,w)\notin R\wedge x\notin Q\right],$$

where Q *is the list of queries from* $\mathcal{A}$ *to* Sim_1*.*

Remark 3.1. It is known [24,43,44,30] that multi-theorem simulation-extractable NIZK (i.e., Definition 3.5) satisfies Definition 3.6; multi-theorem zero-knowledge implies single-theorem zero-knowledge, and simulation-soundness also implies adaptive computational soundness and simulation soundness.

Theorem 3.7 ([44,30]). *Assuming a quantum-secure one-way function and a quantum-secure IND-CPA encryption scheme, any quantum-secure NIZK for NP can be transformed into a quantum-secure simulation-extractable NIZK.*

Corollary 3.1. *Assuming the quantum hardness of LWE, there exists a quantum-secure simulation-extractable NIZK for NP.*

Proof. This follows from Theorems 3.6 and 3.7; IND-CPA encryption can be achieved from the quantum hardness of learning with errors [42]. □

Finally, we define signatures of knowledge.

Definition 3.7 (Signature of Knowledge). *Let R be an* NP *relation with language L and message space $\mathcal{M}$. $\Sigma = \langle \mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify} \rangle$ is a quantum-secure* SimExt*-secure signature of knowledge if the following syntax and properties hold.*

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$: *on input λ, output a* crs *and trapdoor* td.
- $\sigma \leftarrow \mathsf{Sign}(\mathsf{crs}, x, w, m)$: *on input* crs*, pair $(x, w) \in R$, and message $m \in \mathcal{M}$, output a signature of knowledge σ to the message m.*
- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, m, \sigma)$: *on input* crs*, x, m, and σ, accept or reject.*

Correctness: *For every $\lambda \in \mathbb{N}$, pair $(x, w) \in R$ and message $m \in \mathcal{M}$,*

$$\Pr\left[\begin{array}{c}(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) \\ \sigma \leftarrow \mathsf{Sign}(\mathsf{crs}, x, w, m)\end{array} : \mathsf{Verify}(\mathsf{crs}, x, m, \sigma) = 1\right] = 1.$$

Simulation: *There is a* QPT *simulator algorithm $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ such that for every* QPT *algorithm $\mathcal{A}$,*

$$\left|\Pr\left[(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) : \mathcal{A}^{\mathsf{Sign}(\mathsf{crs},\cdot,\cdot,\cdot)}(\mathsf{crs}) = 1\right] - \Pr\left[(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) : \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot,\cdot)}(\mathsf{crs}) = 1\right]\right| \leq \mathsf{negl}(\lambda),$$

where the simulator algorithm Sim_1 discards its fifth input value.

Extraction: *Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulators given by the simulation property. There exist a* QPT *extractor algorithm* Ext *such that for every* QPT *algorithm $\mathcal{A}$*

$$\Pr\left[\begin{array}{c}(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x, m, \sigma) \leftarrow \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot,\cdot)}(\mathsf{crs}) \\ b \leftarrow \mathsf{Verify}(\mathsf{crs}, x, m, \sigma) \\ w \leftarrow \mathsf{Ext}(\mathsf{crs}, \mathsf{td}, x, m, \sigma)\end{array} : \begin{array}{c} b = 1 \ \wedge \ (x, w) \notin R \\ \wedge \ (x, m) \notin Q\end{array}\right] \leq \mathsf{negl}(\lambda).$$

where Q is the list of queries from $\mathcal{A}$ to Sim_1.

Theorem 3.8 ([21,30]). *Given quantum-secure simulation extractable NIZK for NP, there exists a quantum-secure* SimExt*-secure signature of knowledge.*

4 Defining NIZK Arguments with Certified Deletion

In this section, we introduce the notion of a non-interactive zero-knowledge with certified deletion (NIZK-CD) and define its security properties. This primitive allows to transform any NIZK argument for NP to a NIZK construction where proofs can be deleted by the recipients after being verified, and deletion can be substantiated via a deletion certificate. We provide two different definitions of NIZK-CD. The first one is motivated by simplicity, establishes a single-theorem NIZK-CD, and guarantees that no adversary receiving honestly generated proofs for hard NP instances can output both an accepting proof and a valid deletion certificate. In the second definition, we present a multi-theorem NIZK-CD with a guarantee that from any adversary, even having the oracle access to NIZK-CD simulators, returning both an accepting proof and deletion certificate for some instance, one can extract a valid witness corresponding to the instance.

Definition 4.1 (Hard Distribution). *Given an NP relation R, an efficiently samplable distribution $(\mathcal{X}, \mathcal{W})$ over R is hard if for every* QPT *algorithm $\mathcal{A}$*

$$\Pr[(x,w) \leftarrow (\mathcal{X},\mathcal{W}) : (x,\mathcal{A}(x)) \in R] \leq \mathsf{negl}(\lambda).$$

Definition 4.2 (NIZK with Certified Deletion). *Let any* NP *relation R with language L. $\Gamma = \langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}, \mathsf{Delete}, \mathsf{Certify} \rangle$ is a NIZK-CD if it satisfies the following syntax and properties.*

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$*: on input λ, output a classical* crs *and trapdoor* td*.*
- $(\pi, \mathsf{ck}) \leftarrow \mathsf{Prove}(\mathsf{crs}, x, w)$*: on input* crs *and a hard NP instance-witness pair $(x,w) \in \mathcal{R}$, output a quantum proof π and a classical certification key* ck*.*
- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, \pi)$*: on input* crs*, x, and π, output accept or reject.*
- $\mathsf{cert} \leftarrow \mathsf{Delete}(\pi)$*: on input π, output a classical deletion certificate* cert*.*
- $\{0,1\} \leftarrow \mathsf{Certify}(\mathsf{ck}, \mathsf{cert})$*: on input* ck *and* cert*, output accept or reject.*

Perfect Completeness: *For every $\lambda \in \mathbb{N}$ and every $(x,w) \in R$,*

$$\Pr\left[\begin{array}{cc} (\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) & \mathsf{Verify}(\mathsf{crs},x,\pi) = 1 \\ (\pi,\mathsf{ck}) \leftarrow \mathsf{Prove}(\mathsf{crs},x,w) : & \wedge \\ \mathsf{cert} \leftarrow \mathsf{Delete}(\pi) & \mathsf{Certify}(\mathsf{ck},\mathsf{cert}) = 1 \end{array}\right] = 1.$$

Adaptive Computational Soundness: *For every* QPT *algorithm $\mathcal{A}$*

$$\Pr\left[\begin{array}{c} (\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) \\ (x,\pi) \leftarrow \mathcal{A}(\mathsf{crs}) \end{array} : \ \mathsf{Verify}(\mathsf{crs},x,\pi) = 1 \wedge x \notin L\right] \leq \mathsf{negl}(\lambda).$$

Adaptive Computational Zero-Knowledge: *There exists a* QPT *algorithm $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ such that for every* QPT *algorithm $\mathcal{A}$,* QPT *distinguisher $\mathcal{D}$,*

$$\left|\Pr\left[\begin{array}{cc} (\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) & \\ (x,w,\xi) \leftarrow \mathcal{A}(\mathsf{crs}) & : \mathcal{D}(\mathsf{crs},x,\pi,\xi) = 1 \\ (\pi,\mathsf{ck}) \leftarrow \mathsf{Prove}(\mathsf{crs},x,w) & \end{array}\right] - \right.$$

$$\left.\Pr\left[\begin{array}{cc} (\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) & \\ (x,w,\xi) \leftarrow \mathcal{A}(\mathsf{crs}) & : \mathcal{D}(\mathsf{crs},x,\pi,\xi) = 1 \\ (\pi,\mathsf{ck}) \leftarrow \mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x) & \end{array}\right]\right| \leq \mathsf{negl}(\lambda).$$

Deletion Security: *For every* QPT *algorithm* $\mathcal{A}$ *and hard distribution* $(\mathcal{X}, \mathcal{W})$ *over* $\mathcal{R}$*,*

$$\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)\\ (x,w) \leftarrow (\mathcal{X},\mathcal{W})\\ (\pi,\mathsf{ck}) \leftarrow \mathsf{Prove}(\mathsf{crs},x,w)\\ (\pi^*,\mathsf{cert}) \leftarrow \mathcal{A}(\mathsf{crs},x,\pi)\end{array} : \begin{array}{c}\mathsf{Verify}(\mathsf{crs},x,\pi^*)=1\\ \wedge\\ \mathsf{Certify}(\mathsf{ck},\mathsf{cert})=1\end{array}\right] \leq \mathsf{negl}(\lambda).$$

We present the definition of multi-theorem simulation-extractable NIZK-CD.

Definition 4.3 (Simulation-Extractable NIZK-CD). *Let* R *be an* NP *with language* L*.* $\Gamma = \langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}, \mathsf{Delete}, \mathsf{Certify}\rangle$ *is a simulation-extractable NIZK-CD if it satisfies Definition 4.2 and the following properties.*

Adaptive Multi-Theorem Computational Zero-Knowledge: *There is a* QPT *simulator algorithm* $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ *(where* Sim_1 *discards its fourth input) such that for every* QPT *algorithm* $\mathcal{A}$*,*

$$\left|\Pr\left[(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda) : \mathcal{A}^{\mathsf{Prove}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs})=1\right] - \Pr\left[(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) : \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})=1\right]\right| \leq \mathsf{negl}(\lambda),$$

where $\mathcal{A}$ *only receives proofs from the oracles and certifying keys are discarded..*

Simulation Extractability: *Let* $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ *be the simulator given by the adaptive multi-theorem computational zero-knowledge property. There is a* QPT *extractor* Ext *such that for any* QPT *adversary* $\mathcal{A}$*,*

$$\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)\\ (x,\pi) \leftarrow \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})\\ b \leftarrow \mathsf{Verify}(\mathsf{crs},x,\pi)\\ w \leftarrow \mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)\end{array} : \begin{array}{c} b=1 \ \wedge \ (x,w)\notin R\\ \wedge\ x \notin Q\end{array}\right] \leq \mathsf{negl}(\lambda),$$

where Q *is the list of queries from* $\mathcal{A}$ *to* Sim_1*, and* $\mathcal{A}$ *only receives the proof.*

Simulation Extractability with Deletion: *Let* $(\mathsf{Sim}_0, \mathsf{Sim}_1)$ *be the simulator given by the adaptive multi-theorem computational zero-knowledge propery. There exists a* QPT *extractor* Ext-Del *s.t. for every* QPT *algorithm* $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$*,*

$$\Pr\left[\begin{array}{c}(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)\\ (x,\xi) \leftarrow \mathcal{A}_0^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})\\ (\pi,\mathsf{ck}) \leftarrow \mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x,\perp)\\ (\pi^*,\mathsf{cert}) \leftarrow \mathcal{A}_1^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs},\xi,x,\pi)\\ b_1 \leftarrow \mathsf{Certify}(\mathsf{ck},\mathsf{cert}),\ b_2 \leftarrow \mathsf{Verify}(\mathsf{crs},x,\pi^*)\\ w \leftarrow \mathsf{Ext\text{-}Del}(\mathsf{crs},\mathsf{td},x,\pi^*,\mathsf{cert})\end{array} : \begin{array}{c} b_1=b_2=1 \wedge\\ x\notin Q \ \wedge \ (x,w)\notin R\end{array}\right]$$

where Q *is the list of queries, and* $\mathcal{A}$ *only receives the proof from* Sim_1*.*

Remark 4.1. In the definition of simulation extractability with deletion, π^* is a quantum proof output by an adversary and so it may change when operated on by a general quantum algorithm. (Prior work [30,26] was not explicit about this point.) For concreteness, our definition requires π^* to be verified before the extractor tries to extract a witness from the (possibly modified) π^*. We thank an anonymous reviewer from Asiacrypt 2025 for pointing out this subtlety.

The connections between different notions of zero-knowledge, soundness, and extraction for NIZK arguments that are described in Remark 3.1 also hold for NIZK-CD. The following states that the simulation extractability with deletion implies both simulation extractibility and deletion security.

Theorem 4.1. *Simulation extractability with deletion defined in Definition 4.3 implies simulation extractibility as defined in Definition 4.3 and deletion security as defined in Definition 4.2.*

Proof. First, consider an adversary $\mathcal{B}$ that breaks simulation extractibility with a non-negligible advantage. One can build an adversary $\mathcal{A} = (A_0, A_1)$ that breaks simulation extractibility with deletion with the same advantage. In particular, A_0 runs B to get an instance x and a corresponding proof π^*. The instance x is submitted to the simulator Sim_1, and π^* is included in ξ. The algorithm A_1 receives a proof π from Sim_1, which can be deleted to generate a valid deletion certificate cert. Afterwards, A_1 can return produced cert as a valid certificate and π^* included in ξ as an accepting proof corresponding to the instance x.

Similarly, consider an adversary $\mathcal{B}$ that can break the deletion security with a non-negligible advantage. One can build an adversary $\mathcal{A} = (A_0, A_1)$ where A_0 samples a hard instance x and submits it to Sim_1. A_1 receives a proof π corresponding to the instance x from the oracle. Then, A_1 runs B on input x and π to get a valid deletion certificate and an accepting transcript where the success probability is the same as the success probability of $\mathcal{B}$. □

5 Constructions of NIZK-CD from Standard Assumptions

In this section, we propose NIZK-CD constructions from standard cryptographic assumptions. First, we propose a construction with a quantum prover based on quantum-hard one-way functions and standard NIZK proofs. Then, we show how to make the prover classical under the LWE assumption.

5.1 NIZK-CD with a Quantum Prover

Theorem 5.1. *Assuming quantum-hard one-way functions and quantum-secure simulation-extractable NIZK, there exists a simulation-extractable NIZK-CD.*

Proof. Let $\Sigma = \langle \mathsf{Setup}, \mathsf{Com} \rangle$ be a quantum secure, statistically binding and computationally hiding commitment scheme (which can be realized from one-way functions [23]), and $\Pi = \langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify} \rangle$ be any simulation-extractable NIZK. Our NIZK-CD construction is as follows.

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$: $\mathsf{crs}_\Sigma, \mathsf{td}_\Sigma \leftarrow \Sigma.\mathsf{Setup}(1^\lambda)$ and $\mathsf{crs}_\Pi, \mathsf{td}_\Pi \leftarrow \Pi.\mathsf{Setup}(1^\lambda)$, and output the parameters $\mathsf{crs} := (\mathsf{crs}_\Sigma, \mathsf{crs}_\Pi)$ and $\mathsf{td} := (\mathsf{td}_\Sigma, \mathsf{td}_\Pi)$.

- $(\pi, \mathsf{ck}) \leftarrow \mathsf{Prove}(\mathsf{crs}, x, w)$: For $n(\lambda), \ell(\lambda)$, sample randomness $r_{i,b} \leftarrow \{0,1\}^{\ell(\lambda)}$, and $\forall i \in [n], \forall b \in \{0,1\}$ compute $c_{i,b} = \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r_{i,b})$. Then, given the instance $x^* := (x, \{c_{i,b}\}_{i\in[n],b\in\{0,1\}})$ and witness $w^* := (w, \{r_{i,b}\}_{i\in[n],b\in\{0,1\}})$, invoke $\Pi.\mathsf{Prove}(\mathsf{crs}, x^*, w^*)$ and generate a NIZK proof σ for the relation R^*,

$$((x,w) \in R) \quad \bigvee_{i\in[n]} \quad \left(\wedge_{b\in\{0,1\}}\, c_{i,b} = \mathsf{Com}(\mathsf{crs}_\Sigma, (1-b)^{n(\lambda)}, r_{i,b})\right). \tag{1}$$

  Given uniform bits $u_i \leftarrow \{0,1\}$ for all $i \in [n]$, prepare

$$|\psi\rangle := \bigotimes_{i\in[n]} \frac{1}{\sqrt{2}} \left(|0\rangle|r_{i,0}\rangle + (-1)^{u_i}|1\rangle|r_{i,1}\rangle\right). \tag{2}$$

  Output $\pi := (\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle)$ and $\mathsf{ck} := \{u_i, r_{i,b}\}_{i\in[n],b\in\{0,1\}}$.

- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, \pi)$: Parse π as $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle$, run the verifier $v_\sigma \leftarrow \Pi.\mathsf{Verify}(\mathsf{crs}, x^*, \sigma)$, where $x^* := (x, \{c_{i,b}\}_{i\in[n],b\in\{0,1\}})$. Let $U^{\mathsf{com}}_{c_0,c_1}$ be a unitary operation with respect to the commitments c_0 and c_1 as follows.

$$|b\rangle|r\rangle|0\rangle \xrightarrow{U^{\mathsf{com}}_{c_0,c_1}} |b\rangle|r\rangle|\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, b, r, c_b)\rangle, \tag{3}$$

  where the commit-and-compare function is defined as:

$$\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, b, r, c_b) = \begin{cases} 1 & \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r) = c_b \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

  Add ancilla qubits, and apply $U = \bigotimes_{i\in[n]} U^{\mathsf{com}}_{c_{i,0},c_{i,1}}$ to $|\psi\rangle$,

$$\begin{aligned} &\xrightarrow{U} \bigotimes_{i\in[n]} \frac{1}{\sqrt{2}}(|0\rangle|r_{i,0}\rangle|\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, 0, r_{i,0}, c_{i,0})\rangle \\ &+ (-1)^{u_i}|1\rangle|r_{i,1}\rangle|\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, 1, r_{i,1}, c_{i,1})\rangle). \end{aligned} \tag{5}$$

  $\forall i \in [n]$, measure the right-most registers $v_i = \mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, b, r_{i,b}, c_{i,b})$. Output $v = v_\sigma \bigwedge_{i\in[n]} v_i$, where $v = 1$ indicates accept and $v = 0$ otherwise.

- $\mathsf{cert} \leftarrow \mathsf{Delete}(\pi)$: Parse π as $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle = \bigotimes_{i\in[n]}|\psi_i\rangle$ such that $|\psi_i\rangle = \sum_{b\in\{0,1\}} \frac{1}{\sqrt{2}}(-1)^{b\cdot u_i}|b\rangle|r_{i,b}\rangle$. $\forall i \in [n]$, measure each state $|\psi_i\rangle$ in the Hadamard basis, yielding an outcome string $d_i \in \{0,1\}^{\ell(\lambda)}$ such that it holds that $d_i \cdot (r_{i,0} \oplus r_{i,1}) = u_i$. We define the deletion certificate $\mathsf{cert} := \{d_i\}_{i\in[n]}$.

- $\{0,1\} \leftarrow \mathsf{Certify}(\mathsf{ck}, \mathsf{cert})$: Parse ck as $\{u_i, r_{i,b}\}_{i\in[n],b\in\{0,1\}}$ and cert as $\{d_i\}_{i\in[n]}$. Validate whether for all $i \in [n]$ it holds that $d_i \cdot (r_{i,0} \oplus r_{i,1}) = u_i$.

Next, we prove each of the security properties of the proposed construction.

Perfect Completeness: Completeness must be shown with respect to both proof verification and certificate validation. Consider the former case. Following the perfect completeness of Π, we ensure that the proof σ is accepting and $v_\sigma = 1$ with a probability of 1. Then, it suffices to show that for all $i \in [n]$, $v_i = 1$. As the commitment algorithm Com is deterministic, and the commitments are generated honestly, i.e., $c_i = \mathsf{Com}(\mathsf{crs}_\Sigma, b_i^{n(\lambda)}, r_{i,b})$, we ensure that $\forall i \in [n]$ and $\forall b \in \{0,1\}$, we have $\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, b_i, r_{i,b}, c_{i,b}) = 1$. Therefore, the bit v_i is measured as one with a probability of 1. Consider completeness with respect to deletion. We parse the sate $|\psi\rangle$ as $\bigotimes_{i\in[n]}|\psi_i\rangle$, s.t. $|\psi_i\rangle = \sum_{b\in\{0,1\}} \frac{1}{\sqrt{2}}(-1)^{b\cdot u_i}|b\rangle|r_{i,b}\rangle$. The verifier measures each $|\psi_i\rangle$ in the Hadamard basis, which yields an outcome string $d_i \in \{0,1\}^{\ell(\lambda)}$, and it holds that $u_i = d_i \cdot (r_{i,0} \oplus r_{i,1})$. Therefore, the $\mathsf{Certify}$ algorithm always accepts the deletion certificate cert.

Adaptive Multi-Theorem Computational Zero-Knowledge: Consider the simulators $\Pi.Sim = (\Pi.Sim_0, \Pi.\mathsf{Sim}_1)$. We then show $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ for adaptive multi-theorem computational zero-knowledge of our construction.

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)$: Run the algorithms $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi) \leftarrow \Pi.\mathsf{Sim}_1(1^\lambda)$ and $(\mathsf{crs}_\Sigma, \mathsf{td}_\Sigma) \leftarrow \Sigma.\mathsf{Setup}(1^\lambda)$. Output $\mathsf{crs} := (\mathsf{crs}_\Pi, \mathsf{crs}_\Sigma)$ and $\mathsf{td} := (\mathsf{td}_\Pi, \mathsf{td}_\Sigma)$.
- $(\pi, \mathsf{ck}) \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x)$: Sample $r'_{i,b} \leftarrow \{0,1\}^{\ell(\lambda)}$ and $\forall i \in [n], \forall b \in \{0,1\}$ compute $c'_{i,b} = \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r'_{i,b})$. Given $x^* := (x, \{c'_{i,b}\}_{i\in[n],b\in\{0,1\}})$, query $\Pi.\mathsf{Sim}_1$ on input x^* to get σ'. Sample $u'_i \leftarrow \{0,1\}$, prepare $|\psi'\rangle$, and output the proof $\pi := (\{c'_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma', |\psi'\rangle)$ and certificate $\mathsf{ck} := \{u'_i, r'_{i,b}\}_{i\in[n],b\in\{0,1\}}$.

Reduction: Suppose that there exists a QPT adversary algorithm $\mathcal{A}$ such that for some polynomial $p(\lambda)$,

$$\Big|\Pr\Big[\mathsf{crs}, \mathsf{td} \leftarrow \mathsf{Setup}(1^\lambda) : \mathcal{A}^{\mathsf{Prove}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs}) = 1\Big] - \Pr\Big[\mathsf{crs}, \mathsf{td} \leftarrow \mathsf{Sim}_0(1^\lambda) : \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) = 1\Big]\Big| \geq \frac{1}{p(\lambda)}.$$

We construct a QPT adversary $\mathcal{B}$ for the adaptive multi-theorem computational ZK property of the underlying NIZK Π as follows: Receive crs_Π from $\Pi.\mathsf{Setup}$ or $\Pi.\mathsf{Sim}_0$, crs_Σ from $\Sigma.\mathsf{Setup}$, and then send $\mathsf{crs} := (\mathsf{crs}_\Pi, \mathsf{crs}_\Sigma)$ to the adversary $\mathcal{A}$. For each query (x, w), $r_{i,b} \leftarrow \{0,1\}^{\ell(\lambda)}$, $c_{i,b} = \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r_{i,b})$ for all $i \in [n]$ and $b \in \{0,1\}$. Then, given instance $x^* := (x, \{c_{i,b}\}_{i\in[n],b\in\{0,1\}})$ and witness $w^* := (w, \{r_{i,b}\}_{i\in[n],b\in\{0,1\}})$, receive the real or simulated proof σ by query either Prove on input (x^*, w^*) or $\Pi.\mathsf{Sim}_1$ on input x^*. Moreover, sample bits $u_i \leftarrow \{0,1\}$ and prepares $|\psi\rangle$ as defined in Equation 2. The proof $\pi = \{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle$ is sent to the adversary algorithm $\mathcal{A}$. Output the result of $\mathcal{A}$.

As commitment randomness $\{r_{i,b}\}_{i\in[n],b\in\{0,1\}}$ and exponents $\{u_i\}_{i\in[n],b\in\{0,1\}}$ are uniformly sampled, the real and simulated quantum states, i.e., $|\psi\rangle$ and $|\psi'\rangle$, respectively, are statistically indistinguishable. Moreover, computational hiding of the commitment scheme Com ensures that two sets $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}$

and $\{c'_{i,b}\}_{i\in[n],b\in\{0,1\}}$ are computationally indistinguishable. Thus, $\mathcal{B}$ would have a similar and non-negligible polynomial advantage to $\mathcal{A}$ in breaking the adaptive multi-theorem computational zero-knowledge property of Π, i.e., $\frac{1}{p(\lambda)} - \mathsf{negl}(\lambda)$.

Simulation Extractability with Deletion: One can view the commitment algorithm $\mathsf{Com} : \{0,1\}^{\ell(\lambda)+1}$ as one-way function. Assume that there exists an adversary that, on input $\mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r)$, can invert the function and extract b and r with a non-negligible probability. Then, the adversary can break the hiding or binding properties with the same advantage. Lemma 3.4 implies that,

$$\Pr\left[\begin{array}{c} \forall i\in[n] : r_{i,0}, r_{i,1} \leftarrow \{0,1\}^{2\ell(\lambda)}, u_i \leftarrow \{0,1\} \\ \forall i\in[n] : c_{i,0} = \mathsf{Com}(\mathsf{crs}_\Sigma, 0^{n(\lambda)}, r_{i,0}) \wedge c_{i,1} = \mathsf{Com}(\mathsf{crs}_\Sigma, 1^{n(\lambda)}, r_{i,1}) \ : \\ \{(b_i, r_i, d_i)\}_{i\in[n]} \leftarrow \mathcal{A}(\otimes_{i\in[n]} \frac{|0,r_{i,0}\rangle + (-1)^{u_i}|1,r_{i,1}\rangle}{\sqrt{2}}, \{c_{i,b}\}_{i,b}) \\ \wedge_{i\in[n]}\mathsf{Com}(\mathsf{crs}_\Sigma, b_i^{n(\lambda)}, r_i) \in \{c_{i,0}, c_{i,1}\} \\ \wedge \\ \wedge_{i\in[n]} d_i \cdot (r_{i,0} \oplus r_{i,1}) = u_i \end{array}\right] \leq \mathsf{negl}(\lambda)$$

Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the simulators of Π and $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulators of our construction given by the corresponding multi-theorem zero-knowledge. Let $\Pi.\mathsf{Ext}$ be the extractor given by the simulation extractibility of Π. We show an extractor $\mathsf{Ext\text{-}Del}$ that satisfies simulation extractibility with deletion for our proposed construction as follows.

1. On input crs, td, x, π, and cert, parse the proof π as $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle$.
2. Query $\Pi.\mathsf{Ext}$ on input x^* and σ and receive the witness w^*.
3. Output w^* as w.

Reduction: Consider the case that simulation extractibility with deletion does not hold, i.e., there exists a QPT adversary algorithm $A = (A_0, A_1)$ such that given the extractor $\mathsf{Ext\text{-}Del}$, and some polynomial $p(\lambda)$,

$$\Pr\left[\begin{array}{c} (\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x, \xi) \leftarrow \mathcal{A}_0^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) \\ (\pi, \mathsf{ck}) \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x) \\ (\pi^*, \mathsf{cert}) \leftarrow \mathcal{A}_1^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}, \xi, x, \pi) \\ w \leftarrow \mathsf{Ext\text{-}Del}(\mathsf{crs}, \mathsf{td}, x, \pi^*, \mathsf{cert}) \end{array} : \begin{array}{c} \mathsf{Verify}(\mathsf{crs}, x, \pi^*) = 1 \\ \wedge \\ [\mathsf{Certify}(\mathsf{ck}, \mathsf{cert}) = 1 \ \vee \ (x,w) \notin R] \wedge \ x \notin Q \end{array}\right] \geq \frac{1}{p(\lambda)},$$

Since Sim_1 forwards queries to $\Pi.\mathsf{Sim}_1$, we know that x^* is not queried to $\Pi.\mathsf{Sim}_1$. We parse π^* as $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle$. Since $\mathsf{Verify}(\mathsf{crs}, x, \pi^*) = 1$,

we know that $v_\sigma = \Pi.\mathsf{Verify}(\mathsf{crs}_\Pi, x^*, \sigma) = 1$. The witness w^* returned by $\Pi.\mathsf{Ext}$ is a valid witness for R^*. Thus, w^* must include w s.t. $(x, w) \in R$, or, for an $i \in [n]$, it includes $\{r_{i,b}\}_{b\in\{0,1\}}$ s.t. $c_{i,0} = \mathsf{Com}(\mathsf{crs}_\Sigma, 1^{n(\lambda)}, r_{i,0})$ and $c_{i,1} = \mathsf{Com}(\mathsf{crs}_\Sigma, 0^{n(\lambda)}, r_{i,1})$. As $\mathsf{Verify}(\mathsf{crs}, x, \pi^*) = 1$, we can conclude that the measurement outcomes are accepted, i.e., $v_i = 1$ for all $i \in [n]$. According to the definition of U in Equation 5, each commitment $c_{i,b}$ is statistically bound to the message b. Thus, w^* must include a satisfying witness w for R; otherwise, the statistical binding property of the commitment scheme can be broken by running $\mathsf{Ext\text{-}Del}$ to obtain one of the openings and measuring $|\psi\rangle$ in the standard basis to obtain the other opening for the same commitment value.

Then, as $(x, w) \in R$, the adversary needs to output a valid $\mathsf{cert} = \{d_i\}_{i\in[n]}$ to pass the experiment, that is, $\forall i \in [n], d_i \cdot (r_{i,0} \oplus r_{i,1}) = u_i$. In this case, measuring $|\psi\rangle$ in the standard basis yields $\{r_{i,b}\}_{i\in[n]}$ such that $c_{i,0} = \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r_{i,b})$; This contradicts the adaptive hard-core bit property. More precisely, we can build an algorithm $\mathcal{B}$ to break adaptive hard-core bit, where on input $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}$, $|\psi\rangle = \otimes_{i\in[n]} \frac{|r_{i,0}\rangle + (-1)^{u_i}|r_{i,1}\rangle}{\sqrt{2}}$, for an instance x, queries $x^* = (x, \{c_{i,b}\}_{i\in[n],b\in\{0,1\}})$ to $\Pi.\mathsf{Sim}_1$ and obtains a proof σ, generates $\pi := (\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle)$, and queries (x, π) to $\mathcal{A}_1$ to receive an accepting proof π^* and a valid certificate cert. $\mathcal{B}$ can parse π^* as $\{c^*_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma^*, |\psi^*\rangle$ and cert as $\{d_i\}_{i\in[n]}$. Measuring $|\psi^*\rangle$ in the standard basis yields $\{r^*_i\}_{i\in[n]}$ such that $c^*_{i,0} = \mathsf{Com}(\mathsf{crs}_\Sigma, 0^{n(\lambda)}, r^*_i)$ or $c^*_{i,1} = \mathsf{Com}(\mathsf{crs}_\Sigma, 1^{n(\lambda)}, r^*_i)$. $\mathcal{B}$ returns $\{(b_i, r^*_i, d_i)\}_{i\in[n]}$ to pass the experiment. This completes the proof. □

5.2 NIZK-CD with a Classical Prover

In this section, we present our NIZK-CD with a classical prover.

Theorem 5.2. *Assuming quantum hardness of LWE and any quantum-secure simulation-extractable NIZK, there exists a simulation-extractable NIZK-CD s.t. the prover algorithm is classical.*

Proof. Based on the LWE assumption there is a statistically binding, computationally hiding, and computationally equivocal commitment scheme [35] that we denote by $\Sigma = \langle \mathsf{Setup}, \mathsf{Com}\rangle$. Let $\mathcal{F} := \{f_{k,b} : \mathcal{X} \to \mathcal{D}_\mathcal{Y}\}_{k\in\mathcal{K},b\in\{0,1\}}$ be an NTCF family. Let $\Pi = \langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}\rangle$ be a simulation-extractable NIZK. Our NIZK-CD construction with classical prover is described as follows.

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$: $\mathsf{crs}_\Sigma, \mathsf{td}_\Sigma \leftarrow \Sigma.\mathsf{Setup}(1^\lambda)$ and $\mathsf{crs}_\Pi, \mathsf{td}_\Pi \leftarrow \Pi.\mathsf{Setup}(1^\lambda)$, and output the parameters $\mathsf{crs} := (\mathsf{crs}_\Sigma, \mathsf{crs}_\Pi)$ and $\mathsf{td} := (\mathsf{td}_\Sigma, \mathsf{td}_\Pi)$.
- $(\mathsf{ck}, \pi) \leftarrow \mathsf{Prove}(\mathsf{crs}, x, w)$: For $n(\lambda), \ell(\lambda)$, it proceeds in two phases.

  State preparation: The prover and verifier run the following.
  1. The prover runs $\mathsf{NTCF.Gen}_\mathcal{F}$ to generate $\{k_i\}_{i\in[n]}$, $\{\mathsf{td}_{k_i}\}_{i\in[n]}$, and sends the keys to the verifier. The verifier prepares the state

$$|\phi\rangle := \bigotimes_{i\in[n]} \frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x\in\mathcal{X}} |0\rangle|x\rangle + |1\rangle|x\rangle. \tag{6}$$

From Definition 3.2, $|\phi'\rangle$ can be turned into the following superposition.

$$|\phi'\rangle := \bigotimes_{i\in[n]} \frac{1}{\sqrt{2|\mathcal{X}|}} \sum_{x\in\mathcal{X},y\in\mathcal{Y}} \sqrt{(f'_{k_i,0}(x))(y)}|0\rangle|x\rangle|y\rangle + \frac{1}{\sqrt{2|\mathcal{X}|}} \sum_{x\in\mathcal{X},y\in\mathcal{Y}} \sqrt{(f'_{k_i,1}(x))(y)}|1\rangle|x\rangle|y\rangle \tag{7}$$

The verifier measures the images in the standard basis, obtaining the results $\{y_i\}$ and state

$$\bigotimes_{i\in[n]} \frac{1}{\sqrt{2}} \left(|0\rangle|x_{i,0}\rangle + |1\rangle|x_{i,1}\rangle\right). \tag{8}$$

Here, for all $i \in [n]$, the tuple $(x_{i,0}, x_{i,1}) \in \mathcal{X}^2$ is a claw with respect to the measured image $y_i \in \mathcal{Y}$. Finally, the verifier sends the images $\{y_i\}_{i\in[n]}$ to the prover and applies the injection $J : \mathcal{X} \to \{0,1\}^{\ell(\lambda)}$, as defined in Definition 3.2, to the above superposition and prepares the state $|\psi\rangle$ as

$$|\psi\rangle := \bigotimes_{i\in[n]} \frac{1}{2} \left(|0\rangle|r_{i,0}\rangle + |1\rangle|r_{i,1}\rangle\right)$$
$$\forall i \in [n], b \in \{0,1\} : r_{i,b} = J(x_{i,b}). \tag{9}$$

2. For $i \in [n]$, $b \in \{0,1\}$, the prover runs $\mathsf{Inv}(\mathsf{td}_k, b, y_i)$ and $J(x_{i,b})$ to get $r_{i,b}$.

Proof Generation: The remaining parts are similar to the construction in the previous section. The prover generates $c_{i,b} = \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r_{i,b})$ for all $i \in [n], b \in \{0,1\}$. Given $x^* := (x, \{c_{i,b}\})$, $w^* := (w, \{r_{i,b}\})$, the prover runs $\Pi.\mathsf{Prove}(\mathsf{crs}, x^*, w^*)$ and generates a NIZK proof σ for R^*, as defined in Equation 1. It outputs $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma$ and $\mathsf{ck} = \{r_{i,b}\}_{i\in[n],b\in\{0,1\}}$.

In the following, define $\pi = (\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle)$. Note that $|\psi\rangle$ is needed for verification, but the prover never needs access to it.

- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, \pi)$: The verifier runs $v_\sigma \leftarrow \Pi.\mathsf{Verify}(\mathsf{crs}, x^*, \sigma)$. Let $U^{\mathsf{com}}_{c_0,c_1}$ and U be as before (cf. Equation 3). Using U, the state in Equation 5 is prepared. Then, the verifier for all $i \in [n]$, measures the commit-and-compare registers $v_i = \mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, 0, r_{i,0}, c_{i,0})$. The outcome is $v = v_\sigma \bigwedge_{i\in[n]} v_i$.

- $\mathsf{cert} \leftarrow \mathsf{Delete}(\pi)$: The proof π is parsed as $\{c_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma, |\psi\rangle$ and also the state $|\psi\rangle$ as $\bigotimes_{i\in[n]}|\psi_i\rangle$ such that $|\psi_i\rangle = \sum_{b\in\{0,1\}} \frac{1}{\sqrt{2}}|b\rangle|r_{i,b}\rangle$. The verifier then for all $i \in [n]$ measures the state $|\psi_i\rangle$ in the Hadamard basis to get strings $d_i \in \{0,1\}^{\ell(\lambda)}$ such that $d_i \cdot (r_{i,0} \oplus r_{i,1}) = 0$. We define $\mathsf{cert} := \{d_i\}_{i\in[n]}$.

- $\{0,1\} \leftarrow \mathsf{Certify}(\mathsf{ck}, \mathsf{cert})$: Parse ck as $\{r_{i,b}\}_{i\in[n],b\in\{0,1\}}$ and cert as $\{d_i\}_{i\in[n]}$. Validate that for all $i \in [n]$ we have $d_i \cdot (r_{i,0} \oplus r_{i,1}) = 0$ and $d_i \neq \mathbf{0}$.

Then, we prove each of the security properties of our NIZK-CD. Reductions for perfect completeness, adaptive multi-theorem computational zero-knowledge, and simulation extractability with deletion work basically similar to Section 5.1.

Completeness: For proof verification, perfect completeness of Π ensures that $v_\sigma = 1$ with probability 1. As the inversion function Inv and the map J are deterministic, the random coins computed by the prover using trapdoor td_k and the map J are the same as the random coins encoded in $|\psi\rangle$. Moreover, as Com is deterministic and the commitments are generated honestly, we have $\mathsf{Cmt\text{-}Cmp}(\mathsf{crs}_\Sigma, b_i, r_{i,b}, c_{i,b}) = 1$. Thus, for all $i \in [n]$ we have v_i measured as 1. For deletion, we have $|\psi\rangle = \bigotimes_{i\in[n]}|\psi_i\rangle$, s.t. $|\psi_i\rangle = \sum_{b\in\{0,1\}} \frac{1}{2}|b\rangle|r_{i,b}\rangle$. The verifier measures each $|\psi_i\rangle$ in the Hadamard basis to obtain $d_i \in \{0,1\}^{\ell(\lambda)}$ with $d_i \cdot (r_{i,0} \oplus r_{i,1}) = 0$ and $d_i \neq \mathbf{0}$ with overwhelming probability.

Adaptive Multi-Theorem Computational Zero-Knowledge: Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be simulators of Π. Define $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ as:

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)$: Run the algorithms $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi) \leftarrow \Pi.\mathsf{Sim}_1(1^\lambda)$ and $(\mathsf{crs}_\Sigma, \mathsf{td}_\Sigma) \leftarrow \Sigma.\mathsf{Setup}(1^\lambda)$. Output $\mathsf{crs} := (\mathsf{crs}_\Pi, \mathsf{crs}_\Sigma)$ and $\mathsf{td} := (\mathsf{td}_\Pi, \mathsf{td}_\Sigma)$.

- $(\pi, \mathsf{ck}) \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x)$: Run $\mathsf{NTCF.Gen}_\mathcal{F}$, generate $\{k'_i\}_{i\in[n]}$, $\{\mathsf{td}_{k'_i}\}_{i\in[n]}$, and send the key k' to the adversary. For all $i \in [n]$, on input the image $y'_i \in \mathcal{Y}$, compute $x'_{i,0} \leftarrow \mathsf{Inv}(\mathsf{td}_{k'}, 0, y'_i)$ and $x'_{i,1} \leftarrow \mathsf{Inv}(\mathsf{td}_{k'}, 1, y'_i)$, and the randomnesses $r'_{i,0} = J(x'_{i,0})$ and $r'_{i,1} = J(x'_{i,1})$. Compute $\forall i \in [n], \forall b \in \{0,1\}$, $c'_{i,b} = \mathsf{Com}(\mathsf{crs}_\Sigma, b^{n(\lambda)}, r'_{i,b})$. Let $x^* := (x, \{c'_{i,b}\}_{i\in[n],b\in\{0,1\}})$. Query $\Pi.\mathsf{Sim}_1$ on input x^* to get a proof σ'. Prepare $|\psi'\rangle$, as in Equation 9, and output the proof $\pi := (\{c'_{i,b}\}_{i\in[n],b\in\{0,1\}}, \sigma', |\psi'\rangle)$ and the certificate $\mathsf{ck} := \{u'_i, r'_{i,b}\}_{i\in[n],b\in\{0,1\}}$.

  Reduction: It proceeds similarly to the reduction in Section 5.1. The additional note is that the keys k' are uniformly sampled by $\mathsf{NTCF.Gen}_\mathcal{F}$; therefore, they are indistinguishable from the real key k.

Simulation Extractability with Deletion: We can build an extractor $\mathsf{Ext\text{-}Del}$ that satisfies simulation extractibility with deletion and present the reduction the same as Section 5.1. The adaptive hard-core bit property of NTCF hash functions, as defined in Definition 3.2, prevents the adversary from deviating from the protocol. Otherwise, one can use adversary $\mathcal{A}$ to violate at least one of these properties: the statistical binding of the commitment, the simulation extractability of Π, or the adaptive hard-core bit of the NTCF functions. □

6 Applications

In this section, we discuss applications of NIZK-CD to signatures of knowledge and anonymous credentials with certified deletion.

6.1 Signature of Knowledge with Certified Deletion

We formalize and construct the signature of knowledge with certified deletion. This primitive enables the recipient of a signature of knowledge to delete it and generate a deletion certificate. The certificate serves as proof that the recipient no longer retains access to a valid signature.

Definition 6.1 (Signature of Knowledge with Certified Deletion). *For a relation R and a message space $\mathcal{M}$, $\Sigma_R = \langle \mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify}, \mathsf{Delete}, \mathsf{Certify} \rangle$ is a signature of knowledge with certified deletion if it satisfies the following.*

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$: *on input λ, output* crs *and trapdoor* td.
- $(\sigma, \mathsf{ck}) \leftarrow \mathsf{Sign}(\mathsf{crs}, x, w, m)$: *on input* crs, $(x, w) \in R$, $m \in \mathcal{M}$, *output a signature of knowledge σ and a certification key* ck.
- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, m, \sigma)$: *on input* crs, x, m, σ, *accept or reject.*
- $\mathsf{cert} \leftarrow \mathsf{Delete}(\sigma)$*: on input σ, output a deletion certificate* cert.
- $\{0,1\} \leftarrow \mathsf{Certify}(\mathsf{ck}, \mathsf{cert})$*: on input* ck *and* cert, *output accept or reject.*

Correctness: *For every $\lambda \in \mathbb{N}$, pair $(x, w) \in R$ and message $m \in \mathcal{M}$,*

$$\Pr\left[\begin{array}{c} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (\sigma, \mathsf{ck}) \leftarrow \mathsf{Sign}(\mathsf{crs}, x, w, m) \\ \mathsf{cert} \leftarrow \mathsf{Delete}(\sigma) \end{array} : \begin{array}{c} \mathsf{Verify}(\mathsf{crs}, x, m, \sigma) = 1 \\ \wedge \\ \mathsf{Certify}(\mathsf{ck}, \mathsf{cert}) = 1 \end{array}\right] = 1.$$

Simulation: *There is a* QPT *simulator* $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ *such that for every* QPT *algorithm $\mathcal{A}$,*

$$\left| \Pr\left[\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) : \mathcal{A}^{\mathsf{Sign}(\mathsf{crs},\cdot,\cdot,\cdot)}(\mathsf{crs}) = 1\right] - \Pr\left[(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) : \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot,\cdot)}(\mathsf{crs}) = 1\right] \right| \leq \mathsf{negl}(\lambda).$$

where Sim_1 *discards its fifth input value.*

Extraction with Deletion: *Let* $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ *be from the simulation property. There is a* QPT *extractor algorithm* Ext-Del *such that for every* QPT *algorithm $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$,*

$$\Pr\left[\begin{array}{c} (\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x, m) \leftarrow \mathcal{A}_0^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) \\ (\sigma, \mathsf{ck}) \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x, m) \\ (\sigma^*, \mathsf{cert}) \leftarrow \mathcal{A}_1^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}, \sigma) \\ w \leftarrow \mathsf{Ext\text{-}Del}(\mathsf{crs}, \mathsf{td}, x, m, \sigma^*, \mathsf{cert}) \end{array} : \begin{array}{c} \mathsf{Verify}(\mathsf{crs}, x, m, \sigma^*) = 1 \\ \wedge \\ [\mathsf{Certify}(\mathsf{ck}, \mathsf{cert}) = 1 \ \vee \ (x, w) \notin R] \wedge \ (x, m) \notin Q \end{array}\right] \leq \mathsf{negl}(\lambda),$$

where Q is the list of queries from $\mathcal{A}$ to Sim_1.

Remark 6.1. Similarly to NIZK-CD, in our constructions, Setup and $\mathsf{Certify}$ are classical, while Sign, Verify, and Delete are quantum algorithms. Except for σ, which includes quantum information, other parameters are classical. Moreover, Sign might be interactive between a classical signer and a quantum verifier.

Remark 6.2. Extraction with deletion as defined in Definition 6.1 implies extraction as defined in Definition 3.7; this can be proved similarly to Theorem 4.1.

Theorem 6.1. *Assuming any non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge with certified deletion, there exists a signature of knowledge with certified deletion.*

Proof. Let $\Gamma = \langle \mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify}, \mathsf{Delete}, \mathsf{Certify} \rangle$ be a simulation-extractable NIZK-CD. We build a signature of knowledge with certified deletion as below.

- $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)$: Output $\mathsf{crs} \leftarrow \Gamma.\mathsf{Setup}(1^\lambda)$.
- $(\sigma, \mathsf{ck}) \leftarrow \mathsf{Sign}(\mathsf{crs}, x, w, m)$: Let $x^* = (x, m)$ be instance, $w^* = w$ witness for

$$L^* = \{(x, m) : \exists w \text{ s.t. } (x, w) \in R\}.$$

  Then, we have $(\sigma, \mathsf{ck}) \leftarrow \Gamma.\mathsf{Prove}(\mathsf{crs}, x^*, w^*)$ with respect to L^*.
- $\{0,1\} \leftarrow \mathsf{Verify}(\mathsf{crs}, x, m, \sigma)$: Given $x^* = (x, m)$, output $v = \Gamma.\mathsf{Verify}(\mathsf{crs}, x^*, \sigma)$.
- $\mathsf{cert} \leftarrow \mathsf{Delete}(\sigma)$: Output $\mathsf{cert} = \Gamma.\mathsf{Delete}(\sigma)$.
- $\{0,1\} \leftarrow \mathsf{Certify}(\mathsf{ck}, \mathsf{cert})$: Output $v = \Gamma.\mathsf{Certify}(\mathsf{ck}, \mathsf{cert})$.

Next, we prove the correctness, simulation, and extraction with deletion.

Correctness: Since the scheme Γ satisfies perfect completeness, for any honestly generated σ and cert, both satisfy Verify and $\mathsf{Certify}$, respectively.

Simulation: Let $\Gamma.Sim = (\Gamma.Sim_0, \Gamma.\mathsf{Sim}_1)$ be the simulators of Γ. We build $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ for our construction as follows.

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)$: Output $(\mathsf{crs}_\Gamma, \mathsf{td}_\Gamma) \leftarrow \Gamma.\mathsf{Sim}_1(1^\lambda)$.
- $\sigma, \mathsf{ck} \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x, m)$: Output $(\sigma, \mathsf{ck}) \leftarrow \Gamma.\mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x^* := (x, m))$.

Reduction: Suppose a QPT adversary $\mathcal{A}$ exists such that for polynomial $p(\lambda)$, it breaks the simulation property of Σ_R with an advantage greater than $\frac{1}{p(\lambda)}$. We construct a QPT adversary $\mathcal{B}$ for the zero-knowledge property of Γ as follows.

1. Receive real or simulated crs_Γ and send it to the adversary $\mathcal{A}$.
2. For each query (x, m, w), define $x^* := (x, m)$ and witness $w^* := w$, receive the signature σ by a query to Sign or $\Gamma.\mathsf{Sim}_1$, and send σ to $\mathcal{A}$.
3. Output the result of $\mathcal{A}$.

$\mathcal{B}$ has the same advantage $\frac{1}{p(\lambda)}$ in breaking the zero-knowledge property of Γ.

Extraction with deletion: Let $\Gamma.\mathsf{Sim} = (\Gamma.\mathsf{Sim}_0, \Gamma.\mathsf{Sim}_1)$ be the simulators of Γ and $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulators of our construction. Let $\Gamma.\mathsf{Ext}$ be the extractor of Γ. We show an extractor $\mathsf{Ext\text{-}Del}$ for extractibility with deletion.

1. On input crs, td, x, m, σ, and cert, run $\Pi.\mathsf{Ext}(\mathsf{crs}, \mathsf{td}, (x, m), \sigma, \mathsf{cert})$, and receive the witness w^*.
2. Output w^* as w.

Reduction: Let $A = (A_0, A_1)$ a QPT algorithm such that given the extractor $\mathsf{Ext\text{-}Del}$, and a polynomial $p(\lambda)$, it breaks the extraction with deletion property of Σ_R with an advantage greater than $\frac{1}{p(\lambda)}$. Then, one can build an adversary algorithm $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1)$ to break the simulation extractibility with deletion of Γ; $\mathcal{B}_0$ outputs $x^* = (x, m)$ from $\mathcal{A}_0$ and $\mathcal{B}_1$ outputs $(\sigma^*, \mathsf{cert})$ from $\mathcal{A}_1$. □

Corollary 6.1. *Assuming one-way functions and simulation-extractable NIZK, there exists a signature of knowledge with certified deletion. .*

Proof. This follows from Theorem 5.1 and Theorem 6.1. □

Corollary 6.2. *Assuming LWE and simulation-extractable NIZK, there exists a signature of knowledge with certified deletion and classical signing algorithms.*

Proof. This follows from Theorem 5.2 and Theorem 6.1. □

6.2 Anonymous Credentials with Certified Deletion

We construct anonymous credentials with certified deletion [30], where proofs of possession of a credential can be certifiably deleted. We provide only an informal overview of the solution; formal definitions (building on prior work) are fairly straightforward.

The main idea of the construction is as follows. An issuer generates public and private keys for a standard (quantum-secure) signature scheme. A credential on a vector of attributes $\boldsymbol{a} = (a_1, \ldots, a_n)$ is simply a signature on $\boldsymbol{a}$. To prove possession of a credential on revealed attributes $\{a_i\}_{i \in S}$ (while hiding attributes $\{a_i\}_{i \in [n] \setminus S}$), the holder of a credential generates an NIZK-CD that there exists a signature under the issuer's public key on a vector $\boldsymbol{a}$ with $\boldsymbol{a}_i = a_i$ for all $i \in S$.

References

1. Aaronson, S.: Quantum copy protection and quantum money. In: 24th Annual IEEE Conference on Computational Complexity. IEEE (Jul 2009). https://doi.org/10.1109/ccc.2009.42, http://dx.doi.org/10.1109/CCC.2009.42
2. Aaronson, S., Christiano, P.: Quantum money from hidden subspaces. In: 44th Annual ACM Symposium on Theory of Computing (STOC). pp. 41–60. ACM, New York, NY, USA (2012). https://doi.org/10.1145/2213977.2213983, https://doi.org/10.1145/2213977.2213983

3. Aaronson, S., Liu, J., Liu, Q., Zhandry, M., Zhang, R.: New approaches for quantum copy protection. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology—Crypto 2021. pp. 526–555. Springer (2021)
4. Acar, T., Nguyen, L.: Revocation for delegatable anonymous credentials. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) Public Key Cryptography—PKC 2011. pp. 423–440. Springer (2011)
5. Agrawal, S., Kitagawa, F., Nishimaki, R., Yamada, S., Yamakawa, T.: Public key encryption with secure key leasing. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology—Eurocrypt 2023. pp. 581–610. Springer (2023)
6. Amos, R., Georgiou, M., Kiayias, A., Zhandry, M.: One-shot signatures and applications to hybrid quantum/classical authentication. In: 52nd Annual ACM Symposium on Theory of Computing (STOC). p. 255–268. ACM, New York, NY, USA (2020). https://doi.org/10.1145/3357713.3384304, https://doi.org/10.1145/3357713.3384304
7. Ananth, P., La Placa, R.L.: Secure software leasing. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology—Eurocrypt 2021. pp. 501–530. Springer (2021)
8. Ananth, P., Poremba, A., Vaikuntanathan, V.: Revocable cryptography from learning with errors. In: Rothblum, G., Wee, H. (eds.) Theory of Cryptography. pp. 93–122. Springer (2023)
9. Bartusek, J., Goyal, V., Khurana, D., Malavolta, G., Raizes, J., Roberts, B.: Software with certified deletion. In: Joye, M., Leander, G. (eds.) Advances in Cryptology—Eurocrypt 2024. pp. 85–111. Springer (2024)
10. Bartusek, J., Khurana, D.: Cryptography with certified deletion. In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology—Crypto 2023. pp. 192–223. Springer (2023)
11. Bartusek, J., Khurana, D., Malavolta, G., Poremba, A., Walter, M.: Weakening assumptions for publicly-verifiable deletion. In: Rothblum, G., Wee, H. (eds.) Theory of Cryptography. pp. 183–197. Springer (2023)
12. Bartusek, J., Raizes, J.: Secret sharing with certified deletion. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology—Crypto 2024. pp. 184–214. Springer (2024)
13. Belenkiy, M., Camenisch, J., Chase, M., Kohlweiss, M., Lysyanskaya, A., Sha, H.: Randomizable proofs and delegatable anonymous credentials. In: Halevi, S. (ed.) Advances in Cryptology—Crypto 2009. pp. 108–125. Springer (2009)
14. Bennett, C.H., Brassard, G.: Quantum cryptography: Public key distribution and coin tossing. Theoretical Computer Science **560**, 7–11 (2014). https://doi.org/10.1016/j.tcs.2014.05.025, http://dx.doi.org/10.1016/j.tcs.2014.05.025
15. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications. In: 20th Annual ACM Symposium on Theory of Computing (STOC). p. 103–112. ACM, New York, NY, USA (1988). https://doi.org/10.1145/62212.62222, https://doi.org/10.1145/62212.62222
16. Brakerski, Z., Christiano, P., Mahadev, U., Vazirani, U., Vidick, T.: A cryptographic test of quantumness and certifiable randomness from a single quantum device. In: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). pp. 320–331 (2018). https://doi.org/10.1109/FOCS.2018.00038
17. Broadbent, A., Islam, R.: Quantum encryption with certified deletion. In: Pass, R., Pietrzak, K. (eds.) Theory of Cryptography. pp. 92–122. Springer (2020)
18. Broadbent, A., Lord, S.: Uncloneable quantum encryption via oracles. In: Flammia, S.T. (ed.) 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020). Leibniz International Proceedings in Informatics (LIPIcs), vol. 158, pp. 4:1–4:22. Schloss Dagstuhl – Leibniz-Zentrum für

Informatik, Dagstuhl, Germany (2020), https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2020.4
19. Çakan, A., Goyal, V., Raizes, J.: How to delete without a trace: Certified deniability in a quantum world. Cryptology ePrint Archive, Paper 2024/1832 (2024), https://eprint.iacr.org/2024/1832
20. Chardouvelis, O., Goyal, V., Jain, A., Liu, J.: Quantum key leasing for PKE and FHE with classical lessor. In: Advances in Cryptology – Eurocrypt 2025: 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4–8, 2025, Proceedings, Part III. p. 248–277. Springer-Verlag, Berlin, Heidelberg (2025)
21. Chase, M., Lysyanskaya, A.: On signatures of knowledge. In: Dwork, C. (ed.) Advances in Cryptology—Crypto 2006. pp. 78–96. Springer (2006)
22. Coladangelo, A., Majenz, C., Poremba, A.: Quantum copy protection of compute-and-compare programs in the quantum random oracle model (2022), available at https://ia.cr/2020/1194
23. Di Crescenzo, G., Ishai, Y., Ostrovsky, R.: Non-interactive and non-malleable commitment. In: 30th Annual ACM Symposium on Theory of Computing (STOC). p. 141–150. ACM (1998). https://doi.org/10.1145/276698.276722, https://doi.org/10.1145/276698.276722
24. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string. In: 31st Annual Symposium on Foundations of Computer Science. pp. 308–317. IEEE (1990). https://doi.org/10.1109/FSCS.1990.89549
25. Gottesman, D.: Uncloneable encryption. Quantum Information and Computation **3**(6), 581–602 (2003)
26. Goyal, V., Malavolta, G., Raizes, J.: Unclonable commitments and proofs. In: Boyle, E., Mahmoody, M. (eds.) Theory of Cryptography. pp. 193–224. Springer Nature Switzerland, Cham (2025)
27. Helstrom, C.W.: Quantum detection and estimation theory. J. Statist. Phys. **1**, 231–252 (1969). https://doi.org/10.1007/BF01007479
28. Hiroka, T., Morimae, T., Nishimaki, R., Yamakawa, T.: Quantum encryption with certified deletion, revisited: Public key, attribute-based, and classical communication. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology—Asiacrypt 2021. pp. 606–636. Springer (2021)
29. Holevo, A.: Statistical decision theory for quantum systems. Journal of Multivariate Analysis **3**(4), 337–394 (1973). https://doi.org/https://doi.org/10.1016/0047-259X(73)90028-6, https://www.sciencedirect.com/science/article/pii/0047259X73900286
30. Jawale, R., Khurana, D.: Unclonable non-interactive zero-knowledge. In: Advances in Cryptology – Asiacrypt 2024: 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9–13, 2024, Proceedings, Part IX. p. 94–128. Springer-Verlag, Berlin, Heidelberg (2024)
31. Katz, J., Sela, B.: Secret sharing with publicly verifiable deletion. In: Advances in Cryptology—Eurocrypt 2025, Part III. LNCS, vol. 15603, pp. 278–304. Springer (2025)
32. Kitagawa, F., Nishimaki, R., Yamakawa, T.: Secure software leasing from standard assumptions. In: Nissim, K., Waters, B. (eds.) Theory of Cryptography. pp. 31–61. Springer (2021)

33. Kitagawa, F., Nishimaki, R., Yamakawa, T.: Publicly verifiable deletion from minimal assumptions. In: Rothblum, G., Wee, H. (eds.) Theory of Cryptography. pp. 228–245. Springer (2023)
34. Liu, J., Liu, Q., Qian, L., Zhandry, M.: Collusion resistant copy-protection for watermarkable functionalities. In: Kiltz, E., Vaikuntanathan, V. (eds.) Theory of Cryptography. pp. 294–323. Springer (2022)
35. Lombardi, A., Schaeffer, L.: A note on key agreement and non-interactive commitments. Cryptology ePrint Archive, Paper 2019/279 (2019), https://eprint.iacr.org/2019/279, https://eprint.iacr.org/2019/279
36. Morimae, T., Poremba, A., Yamakawa, T.: Revocable quantum digital signatures. In: Magniez, F., Grilo, A.B. (eds.) 19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024). Leibniz International Proceedings in Informatics (LIPIcs), vol. 310, pp. 5:1–5:24. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2024). https://doi.org/10.4230/LIPIcs.TQC.2024.5, https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2024.5
37. Morimae, T., Yamakawa, T.: Quantum commitments and signatures without one-way functions. In: Advances in Cryptology—Crypto 2022, Part I. p. 269–295. Springer (2022). https://doi.org/10.1007/978-3-031-15802-5_10, https://doi.org/10.1007/978-3-031-15802-5_10
38. Peikert, C., Regev, O., Stephens-Davidowitz, N.: Pseudorandomness of ring-LWE for any ring and modulus. In: 49th Annual ACM SIGACT Symposium on Theory of Computing. p. 461–473. STOC 2017, ACM, New York, NY, USA (2017). https://doi.org/10.1145/3055399.3055489, https://doi.org/10.1145/3055399.3055489
39. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology—Crypto 2019. pp. 89–114. Springer (2019)
40. Poremba, A.: Quantum proofs of deletion for learning with errors. In: Tauman Kalai, Y. (ed.) 14th Innovations in Theoretical Computer Science Conference (ITCS 2023). Leibniz International Proceedings in Informatics (LIPIcs), vol. 251, pp. 90:1–90:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2023). https://doi.org/10.4230/LIPIcs.ITCS.2023.90, https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2023.90
41. Radian, R., Sattath, O.: Semi-quantum money. In: 1st ACM Conference on Advances in Financial Technologies (AFT). p. 132–146. ACM, New York, NY, USA (2019). https://doi.org/10.1145/3318041.3355462, https://doi.org/10.1145/3318041.3355462
42. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6) (2009). https://doi.org/10.1145/1568318.1568324, https://doi.org/10.1145/1568318.1568324
43. Sahai, A.: Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In: 40th Annual Symposium on Foundations of Computer Science. pp. 543–553 (1999). https://doi.org/10.1109/SFFCS.1999.814628
44. Santis, A.D., Crescenzo, G.D., Ostrovsky, R., Persiano, G., Sahai, A.: Robust non-interactive zero knowledge. In: Advances in Cryptology—Crypto 2001. p. 566–598. Springer (2001)
45. Unruh, D.: Revocable quantum timed-release encryption. In: Nguyen, P.Q., Oswald, E. (eds.) Adv. in Cryptology—Eurocrypt 2014. pp. 129–146. Springer (2014)
46. Wiesner, S.: Conjugate coding. SIGACT News **15**(1), 78–88 (1983). https://doi.org/10.1145/1008908.1008920, https://doi.org/10.1145/1008908.1008920